

Computer Vision to the Rescue: Infant Postural Symmetry Estimation from Incongruent Annotations

Xiaofei Huang¹, Michael Wan¹, Lingfei Luan¹, Bethany Tunik², Sarah Ostadabbas^{1*}

¹Augmented Cognition Lab (ACLab), Northeastern University, Boston, MA, USA

²Board-Certified Clinical Specialist in Pediatric Physical Therapy

*Corresponding author's email: ostadabbas@ece.neu.edu

Abstract

Bilateral postural symmetry plays a key role as a potential risk marker for autism spectrum disorder (ASD) and as a symptom of congenital muscular torticollis (CMT) in infants, but current methods of assessing symmetry require laborious clinical expert assessments. In this paper, we develop a computer vision based infant symmetry assessment system, leveraging 3D human pose estimation for infants. Evaluation and calibration of our system against ground truth assessments is complicated by our findings from a survey of human ratings of angle and symmetry, that such ratings exhibit low inter-rater reliability. To rectify this, we develop a Bayesian estimator of the ground truth derived from a probabilistic graphical model of fallible human raters. We show that the 3D infant pose estimation model can achieve 68% area under the receiver operating characteristic curve performance in predicting the Bayesian aggregate labels, compared to only 61% from a 2D infant pose estimation model and 60% from a 3D adult pose estimation model, highlighting the importance of 3D poses and infant domain knowledge in assessing infant body symmetry. Our survey analysis also suggests that human ratings are susceptible to higher levels of bias and inconsistency, and hence our final 3D pose-based symmetry assessment system is calibrated but not directly supervised by Bayesian aggregate human ratings, yielding higher levels of consistency and lower levels of inter-limb assessment bias¹.

1. Introduction

Persistent asymmetrical body behavior in early life provides a prominent prodromal risk marker of neurodevelopmental conditions like autism spectrum disorder (ASD), which affects about 2% of children [17, 16, 4, 3]. It is also symptomatic of congenital muscular torticollis (CMT), a

common musculoskeletal condition with an estimated incidence of 3.9% to 16% of infants [8]. Early screening of ASD and CMT is critical for timely intervention and support [14], but currently requires laborious professional behavioral assessments, and for ASD, reliable determinations often only come later in childhood. In this paper, we propose a computer vision method for assessing bilateral infant postural symmetry from images, based on 3D human pose estimation, domain adapted to the challenging setting of infant bodies. Our method appears to be less susceptible to inter-limb biases present in human ratings, and as such could be used to great effect in telehealth, where even experts might find it difficult to judge 3D symmetry from on-screen 2D images. Since our system is based on angles extracted from pose estimation, it is both privacy-preserving and highly interpretable, and is adaptable to new definitions of postural symmetry based on updated scientific hypotheses or discoveries, as well as for different conditions.

Our model assesses bilateral postural asymmetry, first by employing state-of-the-art 3D body pose estimation designed specifically for infant bodies, and second by learning a pose-based assessment calibrated to human ratings of asymmetry. The pipeline is simple but its implementation is highly nontrivial because *reliable ground truth data does not exist for either task*. For pose estimation, there are no infant datasets labeled with 3D ground truth poses, which would require apparatus infeasible for infant subjects. We make some headway by expanding an existing infant body dataset with new 3D pose labels obtained by manual correction of predictions attained from a 3D infant pose estimation model. Nonetheless, as this 3D data is guided only by perception from flat images, these labels can only serve as weak 3D ground truth. As for symmetry assessment, we design a pioneering survey of 10 human raters for their assessments of pose symmetry and angle differences in four pairs of limb across 700 infant images, and find low inter-rater reliability, and suggestions of low internal consistency and high bias. In both settings, ground truth data is constrained by the fundamental challenge of deriving three-dimensional

¹The SyRIP dataset and model code are available at: <https://github.com/ostadabbas/Infant-Postural-Symmetry>

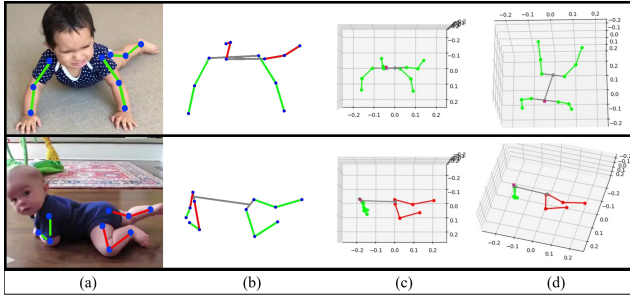


Figure 1: Examples of the discrepancy between 2D pose-based and 3D pose-based symmetry measurements. Left-right pairs of limbs are rendered in green if symmetric and red if asymmetric, according to, respectively, (a) the Bayesian aggregated human ratings of symmetry (i.e. the weak ground truth), (b) 2D pose estimation based assessments, and (c), (d) 3D pose estimation based assessments.

information from two-dimensional images, especially in the domain of infant bodies.

The chief technical thrust in this paper is to “bootstrap” from kernels of reliable information in both tasks to obtain globally reliable and bias-free computer vision assessments of body symmetry. Specifically, our strategy is as follows:

- To remedy the lackluster human rater reliability, we employ a probabilistic graphical model of the human raters as fallible assessors, and compute a Bayesian aggregate of the underlying ground truth, which appears to exhibit a higher level of internal consistency than the human ratings it is derived from. (Full reliability is not assumed.)
- We show that for infant images, the body joint angles obtained from the infant 3D pose estimation model can be used alone to predict the Bayesian ground truth assessments on those images with reasonable accuracy, about 68% area under the receiver operating characteristic (ROC) curve. By comparison, angles obtained from an infant 2D pose estimation model only achieves 61% area under the ROC curve. Some visualized examples of this discrepancy are shown in Fig. 1. The power to predict a response variable obtained completely independently, and despite potential noise in both variables, provides evidence of the accuracy of the 3D pose estimated angles.
- Finally, we learn a simple symmetry classifier for infant images based on the 3D pose estimated angles—now known to be fairly accurate—calibrated by the Bayesian aggregate symmetry rating. This final classifier is *guided* by human intuitions of cutoff thresholds for symmetry assessment, but by design and as we verify quantitatively, is free from the apparent biases stemming from errant factors which affect human judgement. We also demonstrate its superiority over an analogous classifier derived from 2D pose estimates.

Altogether, we offer an exploration of the challenges of human and machine-learning assessment of human body symmetry, and distill our insights into an adaptable, interpretable, end-to-end algorithm for assessing infant symme-

try from still images, with a view towards applications to the early detection and treatment of ASD, CMT, and other common neurodevelopmental conditions.

2. Related Work

Over the past decade, computer vision has been used in the field of automated medical diagnosis, including to distinguish atypical development through video-based behavior monitoring [18, 2]. These vision-based systems provide a low-cost and non-invasive approach, enabling a more objective way to analyze data, and potentially reducing healthcare expenditures when compared to medical examinations. Among key vision-based biomarkers is persistent body asymmetry in infants, which indicates abnormalities associated with developmental disorders, such as ASD and CMT [3, 14].

In [15], as part of the behavioral phenotyping for ASD, authors examined the arm movement and asymmetry in children. They extracted arm and shoulder angles of the child from recorded videos, using a pre-trained real-time multi-person 2D pose estimation model, OpenPose [1]. A computer vision tool to measure and identify ASD behavioral markers based on components of the autism observation was introduced in [5]. Authors first applied 2D pose estimation, which is proposed by extending the Object Cloud Model (OCM) segmentation framework [12] to work with video data, and to produce a 2D stick-man of the toddlers in video segments in which they were walking naturally. Then static and dynamic arm symmetry, as one type of the behavior marker, was detected using the absolute 2D angle difference between corresponding arm parts across time in video segments. Asymmetry was defined if the angle between two corresponding arm parts differs by more than 45° . Both of these papers detected body movement symmetry based on the measured angle differences of arm pairs.

Meanwhile, in [6], authors developed a virtual reality (VR)-based motor intervention methodology by using motion tracking data to quantify efficiency, synchrony and symmetry of whole-body movement. They proposed another kind of hand bilateral symmetry definition, which is the average and standard deviation of the difference in absolute value of horizontal distance between the hands. For symmetry measurement, the 2D locations of wrists were predicted by the pose estimator integrated in the Microsoft Kinect API and then the symmetry score was calculated according to their proposed symmetry measurement formula.

A universal shortcoming of all previous computer vision-based approaches to postural symmetry, however, is their reliance on measurements from 2D body poses, even though human body movement and symmetry is fundamentally three-dimensional. Postural symmetry measurement via 3D body poses has yet to be explored.

3. Concepts and Methods

3.1. Pose-Based Symmetry Measurement

In this paper, we work with a simple parameterized measurement of symmetry for body limbs based on 2D or 3D body joint locations, inspired by the definition of pose symmetry in [3]. First, the infant 2D or 3D pose or skeleton, a collection of human joint locations, is extracted from a flat image by pose estimation algorithms. There are mature computer vision algorithms for this task, but their performance is weaker in the data-scarce infant domain, so we make use of models adapted specifically to infant bodies. For 2D pose extraction, we use the fine-tuned domain-adapted infant pose (FiDIP) model from [7], which works by fine-tuning from an adult pose model to the infant domain, leveraging a domain adversarial network to learn equitably from both real and synthetic infant data. For 3D infant pose detection, a heuristic weakly supervised human pose (HW-HuP) estimation approach [11] is applied. HW-HuP learns partial pose priors from public 3D human pose datasets in flexible modalities, such as RGB, depth or infrared signals, and then iteratively estimates the 3D human pose and shape in the target infant domain in an optimization and regression hybrid cycle. These infant 2D and 3D pose estimators output 17 keypoints and 14 keypoints locations, respectively, but we restrict our poses to the 12 body keypoints needed to define the upper and lower arms and legs (shoulders, elbows, wrists, hips, knees, and ankles), where asymmetry is most prominently manifested.

From the 12 keypoints in the body pose, we can obtain measurements of angles and assessments of symmetry geometrically, as follows, and as illustrated in Fig. 2. First, consider the line segment l_s connecting the two shoulder joints, and then define its mid-perpendicular p_s , the line (in 2D) or plane (in 3D) which intersects l_s orthogonally at its midpoint. Then reflect the upper right arm across p_s , shift it so that its shoulder joint is aligned with that of the left upper arm, and measure the resulting angle. Similarly, reflect the right forearm across p_s , shift it so that its elbow joint is aligned with that of the left forearm, and measure the angle. This is repeated for the legs: reflect, align, and compare the right versus left upper and lower leg angles, this time across the mid-perpendicular of the segment l_h connecting the hip joints. If the formed angle of a given limb pair is less than some fixed predefined angle θ , then the corresponding limb pair is considered to be symmetric, and otherwise it is asymmetric. By adopting above proposed approach and varying angle thresholds, we are able to produce raw angle values and pose symmetry labels for each limb pair in infant images based on their 2D and 3D skeletons.

3.2. Human Symmetry Assessment and Bayesian Aggregation

Pose asymmetry is often assessed by clinical experts to gauge neurodevelopment, or as a symptom for certain de-




Figure 2: The illustration of our infant pose symmetry measurement. The right upper arm and lower arm in red are mirrored across the mid-perpendicular line p_s (in 2D) or plane (in 3D) of the two shoulder joints in green, aligned with their left counterparts at the root joints, resulting in the phantom limbs in blue. The right upper and lower legs are likewise mirrored across the mid-perpendicular p_h of the two hip joints and aligned with their left counterparts. All four resulting angles θ are measured, and the limb pair is considered pose symmetric if its calculated angle is less than a predefined threshold.

velopmental disorders. To guide our algorithmic efforts in emulating clinical evaluations, we commissioned a survey of a number of human raters for their assessments of pose angle and symmetry in infant images, for the pairs of limbs from our symmetry measurement described in Section 3.1. The raters were asked to assess angle differences for limb pairs as per our measurement method, and also to make a subjective judgement of symmetry for each limb pair unguided by this method, to reduce redundancy and to capture information about innate symmetry assessments.

We find that in practise, there is high variation and weak agreement amongst assessments from human raters, and this lack of reliability is not alleviated by simple majority voting, in part because such voting is susceptible to noise from outliers. To rectify this, we employ a probabilistic approach to evaluate different annotators and also give an estimate of the actual hidden labels, as proposed in [13]. When multiple annotators provide possibly noisy labels and there is no absolute gold standard, a maximum-a-posteriori (MAP) estimator is proposed to jointly learn the likely true underlying label and each rater’s intrinsic accuracy. The performance of each rater is measured by calculating sensitivity and specificity with respect to the unknown gold standard, and then a higher weight is assigned to them. We apply an expectation maximization (EM) algorithm to measure the performance of raters based on the given standard, and then optimize the standard based on the new rater performance. The gold standard is initialized by the majority voting result.

Considering that we want to trust some particular raters more than others, a prior is imposed to model the potentially different skill levels of different raters. Beta priors, randomly initialized, are given as conditional information

Table 1: List of data types associated with each infant image considered in this paper. The underlying 700 real infant images are sourced from the synthetic and real infant pose (SyRIP) dataset.

Type	Data	Source
2D body pose	2D coords. of 17 body joints	Included in SyRIP ground truth (Section 4) or inferred from images by 2D pose estimation models (Section 4.2)
3D body pose	3D coords. of 14 body joints	Our weak 3D ground truth labels for SyRIP (Section 4.3), or inferred from images by 3D pose estimation models (Section 4.2)
Pose-based raw angle	Angle in $^\circ$ for 4 limb pairs	<i>Differences</i> between the left and right angles for 4 pairs of limbs (upper arms, lower arms, upper legs, lower legs), derived from 2D or 3D body pose data (above), either ground truth or predicted
Pose-based angle class	3-level angle class for 4 limb pairs	Class from [$<30^\circ$, $30^\circ\text{--}59^\circ$, $\geq 60^\circ$], either inferred from pose based raw angles (above), or canvassed from human raters , (Section 4.1), or voted or Bayesian aggregated from human raters (Section 3.2)
Pose-based symmetry	Binary symmetry class for 4 limb pairs	Either inferred from pose based raw angles (above) based on specified threshold angles, or canvassed from human raters (Section 4.1), or voted or Bayesian aggregated from human raters (Section 3.2)

when calculating the probabilities of sensitivity, specificity, and prevalence for our Bayesian aggregation approach. Following the data gathered in the survey, there are two different types of rating labels: (1) binary class labels, symmetric or asymmetric, and (2) angle class labels [$<30^\circ$, $30^\circ\text{--}59^\circ$, $\geq 60^\circ$], which are intrinsically ordered. For the binary symmetry labels, we infer the likely true human label directly following the EM optimization procedures mentioned above. For the angle class labels, we binarize the variable in one of two different ways—into classes [$<30^\circ$, $\geq 30^\circ$] and into classes [$<60^\circ$, $\geq 60^\circ$]—and then apply the estimation procedure twice, once for each binarization. This yields separate probability estimates for the two binary classes, from which we can infer the most likely class from [$<30^\circ$, $30^\circ\text{--}59^\circ$, $\geq 60^\circ$]. In some rare cases, where the two predicted probabilities are inconsistent with each other, we choose the class with the highest formal probability.

4. Annotation from Humans and Machines

In order to evaluate the performances of human rating and pose-based symmetry measurement, we apply them to a real infant image set of the publicly released synthetic and real infant pose (SyRIP) data [7], which contains 700 real images with assigned posture labels (supine, prone, sitting, and standing) and annotated 2D keypoint locations. See Table 1 for an overview of the data discussed in this paper.

4.1. Human Symmetry Survey

In order to reveal and simulate the mechanism of human rating for postural symmetry, we commissioned survey data from our collaborator, Weifang Zhao of the Chinese National Academy of Arts, who conducted an online experiment study to collect the pose symmetry judgement responses of SyRIP real images from 10 raters through Qualtrics platform. Raters were presented with 700 infant images, divided into 28 sections of 25 images each, with the 25 images within each section randomly shuffled. There were two mandatory 5-minute rest sessions assigned after the 10th and 20th sections. Each image was accompanied

by eight questions: four of them regarding the symmetry of the four limb pairs (upper arm, lower arm, upper leg, and lower leg) and the rest about the predicted angle class between each of the four pairs of limbs. There were five demographic questions at the end of the survey about their major, gender, age, education level, and experience in computer vision or drawing (23 was the mean age; there were 5 male and 5 female participants). A basic snapshot of the survey responses, which plots the mean rater assessment of symmetry at each assessed angle class, can be found in Fig. 9 in the Supplementary Material.

4.2. Infant 2D and 3D Pose Estimation

We employed a range of pose estimation models, listed in Table 2. These include DarkPose [19], a 2D pose estimator trained on large-scale public human pose datasets, and FiDIP [7], which is adapted from DarkPose to work well on infant poses, and trained on SyRIP data. We also used the SPIN [9] 3D pose estimation model, trained on adult 3D motion capture data, and HW-HuP [11], a modification of SPIN designed to predict 3D poses of infants from RGB images, based on training from infant images with depth maps. This depth map data is a weak substitute for the

Table 2: *Top:* Pose estimation models or ground truth data, from which symmetry assessments are derived. *Middle (see Section 5.2):* The optimal threshold angle for each pose estimation model for obtaining symmetry assessments with the highest Cohen’s κ agreement with the Bayesian aggregated symmetry assessments. *Bottom (see Section 5.2):* Areas under the receiver operating curves (Fig. 6) for logistic regression of the Bayesian aggregated symmetry assessment using the raw angles obtained from each pose estimation model.

		Adult Pose Est.	Infant Pose Est.	Ground Truth
Source	2D	DarkPose	FiDIP	SyRIP
	3D	SPIN	HW-HuP	Corrected HW-HuP
Opt. Angle	2D	43.7	42.0	39.1
	3D	17.8	27.9	27.7
ROC AUC	2D	0.60	0.61	0.62
	3D	0.60	0.68	0.73

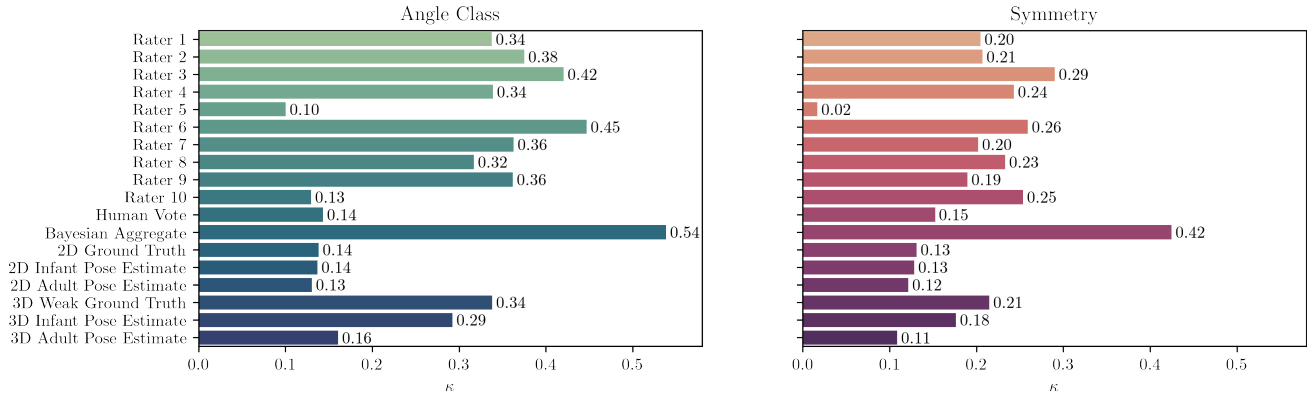


Figure 3: Average Cohen’s κ agreement of a given individual assessment with the 10 human rater assessments (with self-agreement excluded for the human raters). Among human raters, Raters 5 and 10 stand out as outliers for angle assessment, as does Rater 5 for symmetry assessments. The Bayesian aggregate assessment exhibits high average agreement, as expected, but interestingly the human-voted assessment does not. Among the pose-based assessments, those derived from 3D ground truth or 3D predicted poses by using HW-HuP model agree most strongly with human assessments, especially for the more objective assessment of angle level. Since the angle class is ordered, we employ the quadratically weighted Cohen’s κ for those assessments.

3D data available in adult pose datasets, and thus the HW-HuP predictions are not fully reliable. The SyRIP dataset comes with ground truth 2D pose labels, and we next describe how we correct the HW-HuP predictions to obtain weak 3D ground truth labels.

4.3. Infant 3D Pose Correction

To address the lack of reliable 3D pose labels for infants in the SyRIP dataset, we modified the interactive annotation tool introduced in [10] to correct 3D poses predicted by the HW-HuP model introduced above. Since HW-HuP also estimates camera parameters, we overlay its 3D pose keypoint predictions onto the 2D plane over the original infant image, to ensure 2D pose alignment. We interactively modify the global pose orientation and the local bone vector orientation of the 3D skeleton by keyboard inputs to make both the 3D skeleton and the real-time updated projected 2D keypoints locations as correct as possible. In this way, we obtained *weak ground truth 3D pose* labels, limited by inevitable bias from human vision and camera parameter estimation. The distributions of predicted angle differences obtained from various 2D and 3D pose estimation models or ground truth are exhibited in Fig. 10 in the Supplementary Material.

5. Analysis: Computer Vision to the Rescue

We start our analysis by examining shortcomings of human ratings of symmetry, and illustrate how our Bayesian aggregation process ameliorates some of these issues. We then demonstrate the ability of the 3D infant pose estimation models to predict the Bayesian aggregate assessments of angle and symmetry to a higher degree than adult or 2D pose-based models, increasing our confidence in both the Bayesian aggregates and the 3D pose estimations. Finally, we produce our end-to-end symmetry assessments by calibrating the 3D pose-based symmetry assessments with the Bayesian aggregate data. We demonstrate performance

gains afforded by the 3D infant pose-based system over 2D or adult pose-based alternatives, and also illustrate the advantages offered by our algorithmic pose-based assessments over the human and even Bayesian aggregate assessments, both quantitatively and qualitatively. We round out our analysis with two codas on 3D pose estimation and factors affecting symmetry assessment.

5.1. Amending Incongruent Human Annotations

The average Cohen’s κ agreement² between each human rater and their other nine fellow human raters, on their assessments of angle class and symmetry across four pairs of limbs and 700 real images in the SyRIP infant dataset, can be found in Fig. 3. It attests to generally “fair” average agreement for angle class assessments and “slight” to “fair” average agreement for symmetry assessments. In the same vein, the Krippendorff’s α collective agreement amongst the entire group of human raters is 0.30 for angle class and 0.18 for symmetry, attesting respectively to “fair” and “poor” collective agreement. In addition to lower inter-rater agreement, human assessments are also afflicted with high inter-limb assessments agreement for angle class and especially symmetry, as seen in Fig. 4. The high arm-to-leg agreement in symmetry assessments in particular likely indicate unwarranted bias, given that corresponding arm-to-leg agreement for angle class are negligible. Finally, Fig. 5 shows that human ratings of angle class exhibit a low level of correspondence with human ratings of symmetry, suggesting a low level of “internal consistency” among individual human ratings.

Underlying many of these issues is the high variance in assessments between raters (illustrated starkly in the precis of individual rater responses in Fig. 9 in the Supplementary

²Cohen’s κ measures rater agreement on a scale from -1 to 1, while (unlike correlation) accounting for agreements due to random chance. It is often interpreted as a measure of *inter-rater reliability*.

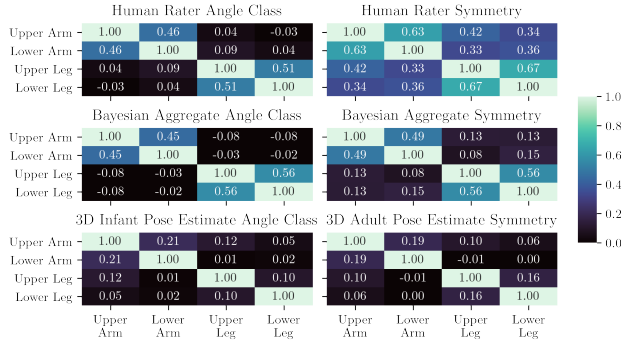


Figure 4: Cohen’s κ agreements between limbs of human, Bayesian aggregate, and 3D pose estimation, and 3D weak ground truth based assessments for angle and symmetry. While some level of inter-limb agreement may exist in the actual (inaccessible) ground truth data, the high agreements in human assessment of symmetry seem particularly excessive, and likely attributable to bias.




Figure 5: Spearman’s ρ ranked correlation between each assessor’s angle and asymmetry assessments across all infants and limb pairs. Assessments with high scores can be interpreted as enjoying high “internal consistency.” Low scores can be caused either by low internal consistency or by angle threshold misalignment (as with the 3D adult pose-based model).

Material), and the high variance of the resulting agreement and consistency metrics. These issues prompt us to explore methods of aggregating the human ratings into a more cohesive whole, including a simple voting method and the probabilistic Bayesian aggregation method described in Section 3.2. The results in Fig. 3, Fig. 4, and Fig. 5 corresponding to these aggregation methods show that the Bayesian aggregate in particular enjoys lower inter-limb agreement and higher angle-asymmetry correspondence than the average human rater—suggesting respectively, lower levels of bias and higher internal consistency—all while maintaining a high level of agreement with the average human rater. Thus, we adopt the Bayesian aggregate assessment as a weak ground truth representation of human assessment of symmetry, with potentially undesirable characteristics excised. Table 3 reports performance metrics of individual human assessments of symmetry, relative to the Bayesian aggregate results as ground truth, again demonstrating the wide variance in human reliability.

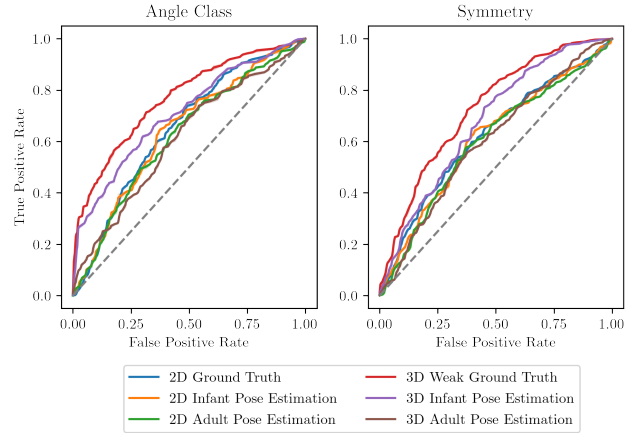


Figure 6: Receiver operating characteristic (ROC) curves for logistic regression of the Bayesian aggregate angle class (left) and symmetry (right), from raw angles from the listed predicted or ground truth sources, across 2800 pairs of limbs in our infant image set. The regression based on our weak 3D ground truth yields best results, while the 3D infant pose estimation model performs better than the other estimation models or the 2D ground truth. The angle classes have been combined into two, [$<30^\circ$, $\geq 30^\circ$], for simplicity. Corresponding areas under the curve (AUCs) can be found in Table 2.

5.2. Pose-Based Symmetry Assessment

As promised, we analyze the extent to which the pose-based systems can track the Bayesian aggregate symmetry assessments, and then calibrate and evaluate our final automated symmetry assessment system. We first consider the raw angle data obtained from pose estimation, which, as described in Section 3.1, consists of a set of four angle differences in degrees, for the four key pairs of limbs under consideration (upper arms, lower arms, upper legs, and lower legs). We gauge agreement with the Bayesian aggregate rater assessments of angle class and of symmetry by logistically regressing for them using the raw angles, for each of the four joints and each infant image (2800 data points in all). Fig. 6 shows the receiver operating characteristic (ROC) curves resulting from this regression, performed with a 3:1 train-test split. For the regression of angle class, we convert the three true classes into a binary variable indicating whether the angle is over 30° for ease of interpretation. The areas under the curve (AUC) for the ROC curves in the symmetry regression are provided in Table 2.

These metrics confirm that the raw angles from the weak 3D ground truth can model the Bayesian aggregate assessment of both angle and symmetry to a high degree of fidelity. Neither set of data can be taken as fully reliable ground truth, but since they are derived from different hu-

Table 3: Performance of each human rater assessment of symmetry relative to Bayesian aggregate.

Rater #	1	2	3	4	5	6	7	8	9	10
Sensitivity	0.98	0.97	0.88	0.54	0.99	0.61	0.99	0.64	0.38	0.79
Specificity	0.35	0.38	0.80	0.96	0.03	0.96	0.30	0.89	0.98	0.80

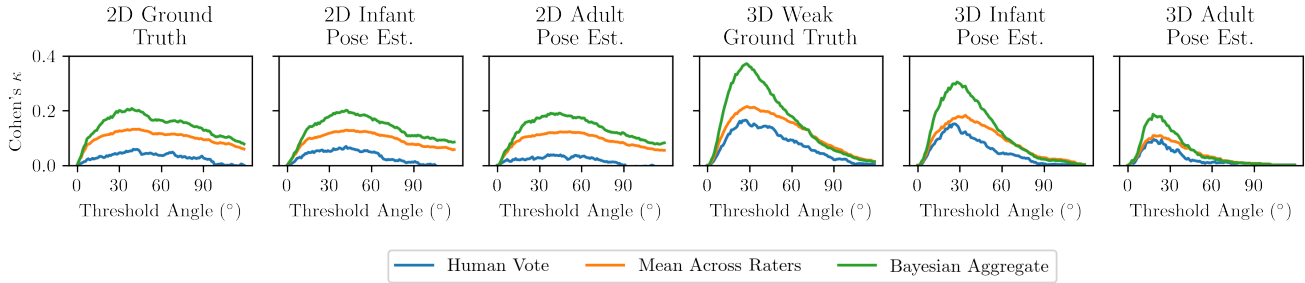


Figure 7: Cohen’s κ agreement between various pose estimation based assessments of symmetry and human assessments, as the threshold angle for defining the threshold for pose base assessment varies. Agreement with human assessment is either given as the mean of agreement with each of the 10 raters, or as agreement with the voted or Bayesian aggregate rater. These results show that the highest capacity for agreement is afforded by the Bayesian aggregate rater on the one hand, and the 3D infant pose estimation based model on the other.




Figure 8: A special example demonstrating the effect of angle threshold on pose-based measurement results and constraints on the feasibility of 2D pose-based method. Left-right pairs of limbs are rendered in green if symmetric and red if asymmetric, according to, respectively, (a) the Bayesian aggregated human ratings of symmetry (i.e. the weak ground truth), (b) 2D pose-based assessments, and (c), (d) 3D pose-based assessments.

man annotators performing fairly different tasks, the high level of agreement exhibited here increases our confidence in the accuracy of both. Among pose estimation models, the 3D infant-specific models enable the next best predictions of human symmetry assessments, while the poses from the remaining models—either 3D pose models for the general most adult human, or 2D pose models for infants or adults—offer weaker ability to predict the human assessments.

We now turn to the task of calibrating an end-to-end pose-based system for the evaluation of symmetry, for use in practical applications or further research³. In concrete terms, wish to select threshold angles which will allow us to convert our raw joint angles to binary assessments of symmetry per joint, in a way that maximizes concordance with the Bayesian aggregate. We choose to guide this concordance with the same Cohen’s κ agreement score employed earlier. Fig. 7 shows the Cohen’s κ agreement of symmetry assessments derived from all six of the pose-based estimators at various decision angle thresholds, compared with both the voted and Bayesian raters; it also shows the mean Cohen’s κ agreement with each of the ten human raters individually. Incidentally, these results not only confirm again the supremacy of the 3D ground truth and infant pose prediction methods for tracking human assessments of symmetry, but on the flip side, also demonstrate the superior-

³For some applications, the output of the four raw angles will suffice as a multi-dimensional, continuous measure of overall body symmetry, affording the end-user the flexibility to redefine the overall notion of symmetry for different tasks or in response to advances in scientific understanding.

ity of the Bayesian aggregation of human symmetry assessments over the voted or the average human assessment for tracking the 3D weak ground truth assessment, at most reasonable angle thresholds. From these Cohen’s κ curves, we extract the threshold angles maximizing agreement for each pose-based model, reported in Table 2. These thresholds then define the corresponding symmetry assessment for each model (or ground truth data).

Metrics quantifying these final assessment models have already been reported throughout the paper, but we offer our interpretations here. Fig. 3 confirms, as expected, that the 3D infant pose estimation assessment offers the highest average agreement with human rater assessment, compared to the 3D adult pose estimation or the 2D infant pose estimation models; our model also comes close to the level achieved by the 3D weak ground truth assessment. Fig. 4 shows that assessments based on predicted or weak ground truth 3D poses are relatively free from inter-limb agreement, compared to individual or aggregate human assessments. In the absence of fully reliable ground truth assessments, this circumstantially suggests that human assessments are susceptible to bias from nearby parts, while our automated approach is not. Finally, Fig. 5 shows that most of the pose estimation based models enjoy high internal consistency in their assessment of angle class versus symmetry, as to be expected from mechanistic models.

5.3. Qualitative Evaluation

We illustrate the performance of the pose-based models in Fig. 1 and Fig. 8. These figures visualize the Bayesian aggregate assessments of symmetry on top of the original image, as well as the assessments derived from the 2D and 3D infant pose models (FiDIP and HW-HuP, respectively) on top of their respective predicted pose skeletons, with green indicating symmetric judgements and red indicating asymmetric judgements.

In Fig. 1, we see examples of infant poses where the 2D pose-based assessment is mistaken, but the 3D pose-based assessment is able to make the correct call, according to the Bayesian aggregate label. We have highlighted multiple

views of the 3D skeleton to highlight the advantage that the 3D pose-based assessment has, and suggest that the mistakes made by the 2D pose-based assessment are in a way understandable, given that it is limited to a single perspective, so to speak. Fig. 8 shows special case where the 2D perspective is actually not too limiting, since the infant is lying flat on its back, so its limbs largely confined to a plane parallel to the image plane. Indeed, in this case, the 2D and 3D pose-based symmetry assessments agree, but they differ from the Bayesian assessment, which may reflect human bias or simply a stricter subjective threshold for symmetry on the part of the human assessments. More comparisons between 2D or 3D pose-based assessments are exhibited in Fig. 11 in the Supplementary Material.

5.4. Coda 1: Improving 3D Pose Estimation

The main factor limiting performance of state-of-the-art infant 3D pose estimators such as HW-HuP is the scarcity of “true” ground truth 3D pose data. We briefly report on the effect of fine-tuning HW-HuP with the weak 3D ground truth labels for SyRIP images generated for this paper. We split the 700 SyRIP images into a 100 image test set, coinciding with SyRIP’s Test100 set, and a 600 image train set, and fine-tune the infant HW-HuP model on the 600 image train set with weak 3D labels, for 200 epochs with learning rate of $5e - 05$. The resulting performance of the fine-tuned infant HW-HuP model under the mean per joint position error (MPJPE), as reported in Table 4, is significantly improved over the base infant HW-HuP model, and over the adult pose SPIN model.

Table 4: 3D pose estimation performance in mean per joint position error (MPJPE) in mm on the 100 image SyRIP test set with our weak ground truth 3D labels. We compare our fine-tuned HW-HuP model (HW-HuP-FT) with the base HW-HuP model, and with the adult-trained SPIN model.

Model	SPIN	HW-HuP	HW-HuP-FT
MPJPE	105.8	97.2	78.3

5.5. Coda 2: Factors Affecting Symmetry Assessments

We conclude our study with a supplementary importance analysis of factors affecting the assessments of symmetry considered in our work, via logistic regression. We take the Bayesian aggregate symmetry assessment as the response variable, and the following as covariate factors: limb part under consideration (upper arm, lower arm, upper leg, or lower leg), the infant posture (included in SyRIP), an occlusion label for each limb (which we annotate for this purpose), and finally, an angle variable. We consider two separate sources for the angle variable, one consisting of the angle class assessments from all of the human raters, and one obtained from the 3D infant pose estimation. We find that both of regressed models are statistically significant.

According to the logistic regression result for Bayesian aggregation, all four predictors account for between 40.8% (R_{CS}^2) and 54.6% (R_N^2) of the variance in the dependent variable and correctly classify 83.2% of cases. From the logistic regression result for four predictors, we conclude that only limb part and estimated angle between corresponding limb part significantly contributed to the asymmetry assessment model. While for 3D prediction model, the logistic regression result indicated that all factors expect occlusion significantly contribute to the model. All four factors explain for between 45.5% (R_{CS}^2) and 60.5% (R_N^2) of the variance in the dependent variable and correctly classify 90.8% of cases.

In order to assess predictor importance, we use the decrease of R_{CS}^2 approach to calculate the ΔR_{CS}^2 when removing one of the predictor. A larger decrease indicates more contribution of the removed predictor to explain the model. The results of decreasing R_{CS}^2 are reported in Table 5. When the angle estimation was taken out of the model, the R_{CS}^2 value declined by -0.342 for the Bayesian model and by -0.366 for the 3D prediction model correspondingly. Thus, angle estimate is the most important predictor of all the variables that have been found, and it is used in both the human rating model and the 3D prediction model.

Table 5: Parameter importance evaluation for logistic regression of Bayesian aggregate symmetry class from human rater and 3D infant pose estimate assessments.

Excluded Feature	Human Raters		3D Infant Pose Est.	
	R_{CS}^2	ΔR_{CS}^2	R_{CS}^2	ΔR_{CS}^2
None (Full Model)	0.408		0.455	
Limb Part	0.402	-0.006	0.418	-0.037
Posture	0.408	0.000	0.448	-0.007
Angle	0.066	-0.342	0.089	-0.366
Occlusion	0.408	0.000	0.455	0.000

6. Conclusion

We have presented a computer vision based method for assessment of postural symmetry in infants from their images, with the goal of enabling early detection and timely treatment of issues related to infant motor and neural development. We found human ratings of symmetry to be unreliable, and rectified them with a Bayesian-based probabilistic aggregate rating. We demonstrated that automatic assessments based on pose estimation avoid some of the pitfalls of human assessments, while retaining the ability to predict the Bayesian aggregate ratings to a strong degree, with 3D infant pose models performing stronger than 2D models or adult pose models.

7. Acknowledgements

This study was supported by the National Science Foundation (NSF-CAREER 2143882). The authors would also like to thank Dr. Weifang Zhao of the Chinese National Academy of Arts for his contributions in data collection.

References

- [1] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017.
- [2] Ryan Anthony J de Belen, Tomasz Bednarz, Arcot Sowmya, and Dennis Del Favero. Computer vision in autism spectrum disorder research: a systematic review of published studies from 2009 to 2019. *Translational psychiatry*, 10(1):1–20, 2020.
- [3] Gianluca Esposito and Paola Venuti. Symmetry in infancy: analysis of motor development in autism spectrum disorders. *Symmetry*, 1(2):215–225, 2009.
- [4] Gianluca Esposito, Paola Venuti, Sandra Maestro, and Filippo Muratori. An exploration of symmetry in early autism spectrum disorders: analysis of lying. *Brain and Development*, 31(2):131–138, 2009.
- [5] Jordan Hashemi, Thiago Vallin Spina, Mariano Tepper, Amy Esler, Vassilios Morellas, Nikolaos Papanikolopoulos, and Guillermo Sapiro. A computer vision approach for the assessment of autism-related behavioral markers. In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, pages 1–7. IEEE, 2012.
- [6] Darren R Hocking, Adel Ardalan, Hisham M Abu-Rayya, Hassan Farhat, Anna Andoni, Rhoshel Lenroot, and Stan Kachnowski. Feasibility of a virtual reality-based exercise intervention and low-cost motion tracking method for estimation of motor proficiency in youth with autism spectrum disorder. *Journal of NeuroEngineering and Rehabilitation*, 19(1):1–13, 2022.
- [7] Xiaofei Huang, Nihang Fu, Shuangjun Liu, and Sarah Ostadabbas. Invariant representation learning for infant pose estimation with small data. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–8. IEEE, 2021.
- [8] Sandra L. Kaplan, Colleen Coulter, and Barbara Sargent. Physical Therapy Management of Congenital Muscular Torticollis: A 2018 Evidence-Based Clinical Practice Guideline From the APTA Academy of Pediatric Physical Therapy. *Pediatric Physical Therapy*, 30(4):240–290, Oct. 2018.
- [9] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2252–2261, 2019.
- [10] Shichao Li, Lei Ke, Kevin Pratama, Yu-Wing Tai, Chi-Keung Tang, and Kwang-Ting Cheng. Cascaded deep monocular 3d human pose estimation with evolutionary training data. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [11] Shuangjun Liu, Xiaofei Huang, Nihang Fu, and Sarah Ostadabbas. Heuristic weakly supervised 3d human pose estimation in novel contexts without any 3d pose ground truth. *arXiv preprint arXiv:2105.10996*, 2021.
- [12] Paulo AV Miranda, Alexandre X Falcão, and Jayaram K Udupa. Cloud models: Their construction and employment in automatic mri segmentation of the brain. 2010.
- [13] Vikas C Raykar, Shipeng Yu, Linda H Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. *Journal of machine learning research*, 11(4), 2010.
- [14] Barbara Sargent, Sandra L Kaplan, Colleen Coulter, and Cynthia Baker. Congenital muscular torticollis: bridging the gap between research and clinical practice. *Pediatrics*, 144(2), 2019.
- [15] JA Sarmiento and PC Naval. Behavioral phenotyping for autism spectrum disorder biomarkers using computer vision. *Journal of Image and Graphics*, 8(2):47–52, 2020.
- [16] Osnat Teitelbaum, Tom Benton, Prithvi K. Shah, Andrea Prince, Joseph L. Kelly, and Philip Teitelbaum. Eshkol-wachman movement notation in diagnosis: The early detection of asperger’s syndrome. *Proceedings of the National Academy of Sciences*, 101(32):11909–11914, 2004.
- [17] Philip Teitelbaum, Osnat Teitelbaum, Jennifer Nye, Joshua Fryman, and Ralph G Maurer. Movement analysis in infancy may be useful for early diagnosis of autism. *Proceedings of the National Academy of Sciences*, 95(23):13982–13987, 1998.
- [18] Jérôme Thevenot, Miguel Bordallo López, and Abdenour Hadid. A survey on computer vision for assistive medical diagnosis from faces. *IEEE journal of biomedical and health informatics*, 22(5):1497–1511, 2017.
- [19] Feng Zhang, Xiatian Zhu, Hanbin Dai, Mao Ye, and Ce Zhu. Distribution-aware coordinate representation for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7093–7102, 2020.